

HP
Planning
and
Architecture Guide
for
Windchill Release 6.2.6

Version 5
July 18, 2003

Michael Finocchiaro
Stephen Fleischman
HPTC ISV Solutions Engineering Group

hp	Error! Bookmark not defined.
Planning	1
and.....	1
Architecture Guide.....	1
for	1
Windchill Release 6.2.6.....	1
Introduction.....	3
Choosing an Architecture	3
3-T or not 3T	3
Vault or Blobs	6
Oracle 32- or 64-bit	6
Oracle 8i or 9i.....	7
Web Servers and Servlet Engines	7
Network Bandwidth	7
Partitions on hp high-end Servers	7
About Partitions	7
Hardware Partitioned Systems	8
Partitioning and Multiple Windchill instances	8
Choosing Disks	9
Configuring for High Availability	9
Mirror/UX	9
MC/ServiceGuard.....	10
Configuring for Security	10
VirtualVault	Error! Bookmark not defined.
Firewalling.....	11
Planning a Backup Strategy	11
Omniback/Data Protector	11
Oracle Archiving and Hot Backup	11

Introduction

Besides installation (cf. Windchill R6.2.6 Installation Guide for HP), there are many other considerations to deploying Windchill in a customer environment. This white paper is intended to help explore some ideas and guidelines of how to make this task manageable and efficient.

Choosing an Architecture

3-T or not 3T

Modern web applications are usually deployed on a multi-layer or 3-Tier architecture. What this means is that the application layer (Web and Java servers) are not placed on the same machine as the backend or database layer. In the case of Windchill on HP-UX, there are two possibilities:

- Monolithic configuration: all the components on the same machine – or –
- 3-Tier configuration with Windchill, Tomcat, and Apache on one machine and Oracle 8i on another machine.

The following table gives some basic differences between these choices:

Feature	3-Tier Architecture	Monolithic Architecture
Scalability	Full CPU's available to Windchill; Very scalable	CPU's shared between Windchill and Oracle; Less scalable
Investment	Requires at least two machines	Requires only one machine
Vaulting	Vault must stay on Windchill machine so disk space must be available to both machines	Requires enough disk for both database and vault
Security	Protection of the database as it is isolated on a different machine	Database vulnerable to system failure
High Availability	Can be integrated in an MC/ServiceGuard cluster	No HA possible
Performance	Measured 30-50% better than monolithic depending on workload	Less performance due to sharing of all CPU and memory resources by Oracle and Windchill
Robustness	Having the two systems separated gives more resource to the front-end and thus the 8:30AM syndrome can be more easily handled	No built-in protection against a sudden burst of activity on Windchill
Manageability	Increased amount of work due to two systems	Easy to manage

Table 1: Comparison of 3-Tier and Monolithic Architectures



Please note that in sizing systems in a highly available configuration using MC/Service Guard, you should plan for the worst-case fail over scenario, i.e., size the Windchill system based on the monolithic sizing data and not on the 3-tier data. This is because if a system failover occurs and if you have sized for the monolithic case, the impact on the users will be as little as possible and the performance should still be sustainable.

As a general rule, as the number of users increases, you will get better performance from a 3-Tier architecture. Also the benefits of High Availability and Security need to be considered, please see the High Availability section later in this report. For a small development site, a monolithic configuration is probably adequate. For a medium-to-large deployment, a 3-Tier is a better answer.

Figure 1 shows an example of a 3-Tier configuration

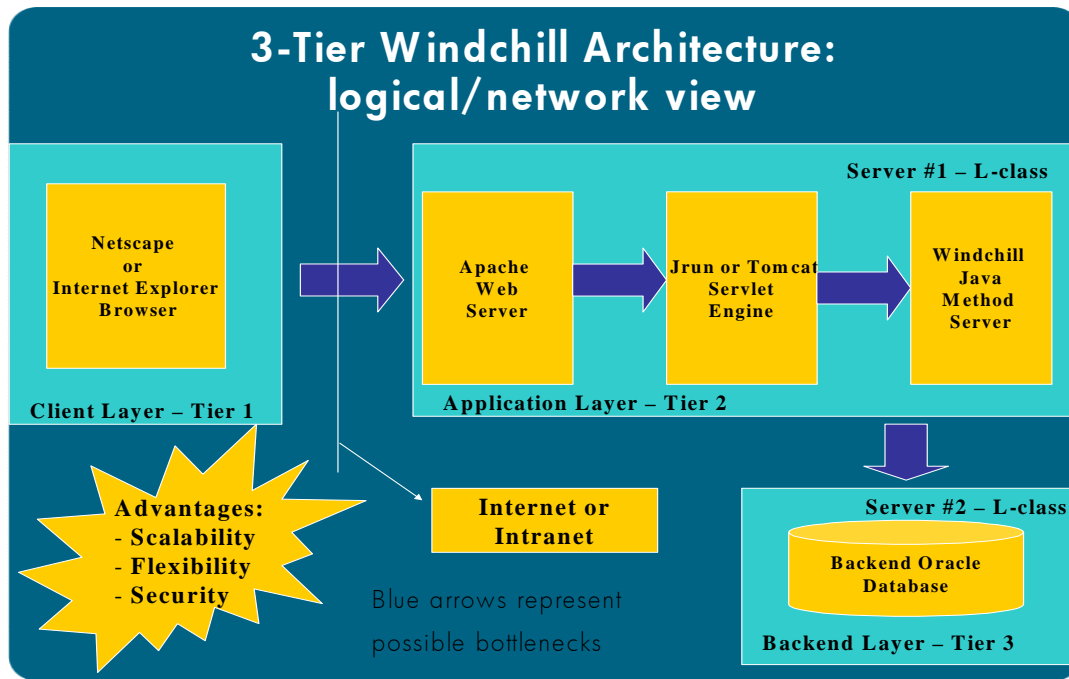


Figure 1. Windchill 3-Tier architecture

Along these same lines, another option to consider is clustering. This means using two or more machines at the application layer and a load balancing mechanism. Windchill has some built-in load balancing capabilities in the Method Server. The additional hardware required is an IP load balancing switch such as the Cisco Local Director for balancing and distributing the incoming traffic. At this time, we do not have sizing numbers available specific to this architecture but it is expected that the 350 users/CPU recommendation will still hold if not improve. See Figure 2 below for an example.

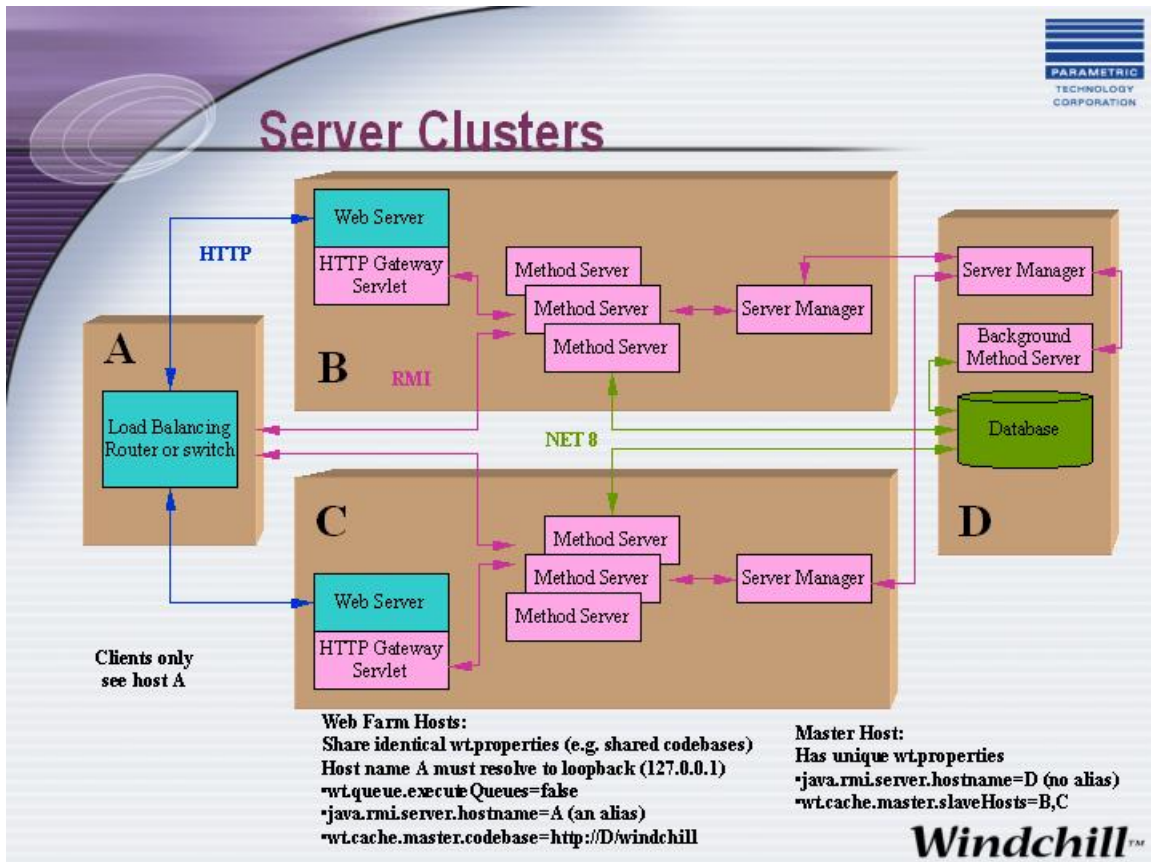


Figure 2. Windchill Server Clusters

For other applications in the Windchill solution stack (RetrievalWare, Aphelion, Workgroup Managers), more servers may be advisable.

- RetrievalWare requires quite a bit of resources and thus is suggested to run on the Oracle database server or on its own 2-processor machine
- Aphelion 2001 Directory Server is relatively light-weight but if there are several projects to manage on separate machines, it is suggested to place it on a separate server. It is not recommended to rely on Aphelion for a global corporate LDAP solution. Typically, iPlanet Directory Server or another commercial product will be used. In this case, there needs to be a connection to the corporate LDAP from the Apache/Windchill server. Only Info*Engine-specific information will be stored in Aphelion and Windchill will be configured to use the corporate LDAP for authentication.
- The Workgroup Managers are large consumers of resources and can require a graphics card. Typically, several UNIX or Wintel workstations are installed on the same network as the Windchill server and CAD agents run for each CADD application to be processed (e.g. CADDs, Pro/E, CATIA, etc). If you are merely doing thumbnails, this can be done on the Windchill server using the Xvfb as described in the Windchill R6.2.6 install guide.
- Additional Info*Engine adapters can also be run either on the Windchill server or on separate servers. If using TIBCO or another EAI (Enterprise Application Integration) package, it is recommended to run this on a separate Wintel machine.

Here is an example of a configuration including some of the elements discussed above:



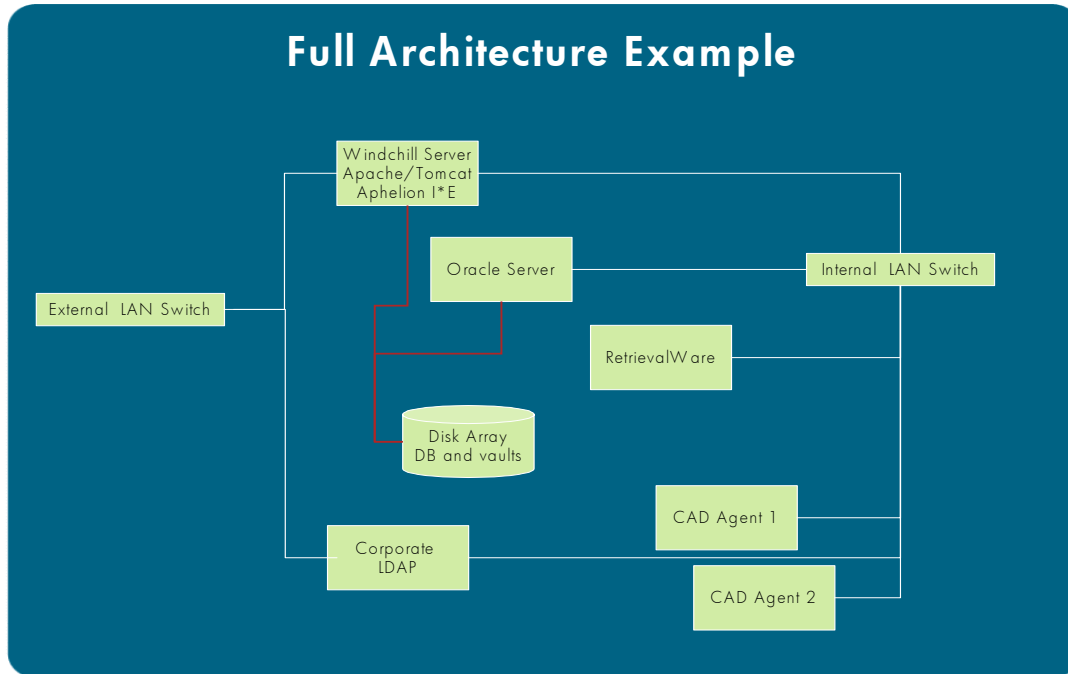


Figure 3 - Full Architecture Example

Vault or Blobs

Windchill offers the option of storing content (CAD files, documents, etc) into either the Oracle database using Binary Large Object Blocks (or BLOBS) or onto an external file system (or vault) with only metadata stored in Oracle.

Primarily this choice depends on the type of data that is being stored. If the data is primarily documents, (e.g., PowerPoint and Word documents, etc), from an administrative point of view, Blobs are probably a better solution. However, if you are storing large CAD data files, a file vault will probably be more ideal. Things to keep in mind for a vaulted configuration:

1. Always locate the vault as close if possible to the Windchill server if not on the Windchill server itself
2. Remember to include the vault in your backup procedure, cf. the Backup section later this report.
3. Think about Windchill File Replication for getting the data as close as possible to its consumers, cf. the Windchill Administrator's Guide.

Note that benchmarking with 1 Mb files did not show an appreciable difference between Blobs and Vault but this is not a realistic model size.

Oracle 32- or 64-bit

Oracle8i V3 (8.1.7) and Oracle 9iR2 (9.2.0.x) on HP-UX comes in two flavors: 32- and 64-bit. The choice between them essentially depends on the size of the database and the available

memory on the Oracle server machine. If the database size is going to be in the 100's of gigabytes up into the terabyte range, you will need to use the 64-bit version for addressability. If your database is greater than 2Gb, the 64-bit version is interesting in that an SGA of greater than 2 Gb can be created to keep as much data as possible in memory. Lastly, if you have a relatively small memory footprint on your Oracle server (less than 4 Gb), then you should stay at the 32-bit version.

Oracle 8i or 9i

Oracle 9i is supported with Windchill / PDMLink / ProjectLink R6.2.6 and above. The benefits of 9i vs. 8i are:

- Cost : you can use an Oracle Standard Edition license for 9i with Windchill as opposed to the Enterprise licenses required when using 8i
- Performance : the JDBC layer is faster and more scalable on 9i than 8i
- Availability : 9i provides RAC (Rapid Application Clusters) functionality which allows clustering of the database over several machines for both scalability and data availability and data security.

Web Servers and Servlet Engines

HP-UX11i comes with Apache and Tomcat pre-installed. This is the recommended configuration as it is the preferred solution from PTC, requires a minimum of setup and is free. Alternatively, iPlanet can be used with its internal servlet engine. JRun is no longer recommended as there are major issues with its support from PTC and from Allaire.

Apache can be run on separate servers from the Windchill server if desired. This would allow for less expensive web-farms front-ending a larger Windchill server.

Network Bandwidth

The scalability testing demonstrated that at 1000 users, the LAN traffic was 600Kb/second that is appropriate for a 10Base-T network. In a 3-tier architecture, it is recommended that the Oracle and Windchill servers be connected into the same 100Mbit switch for better reliability and throughput, but clients on a 10Mbit network are perfectly acceptable. The 2000 user run demonstrated in excess of 1Mb/second of network traffic indicating this would require a 100Mbit network.

Further WAN testing is ongoing at this time.

Partitions on hp high-end Servers

About Partitions

Partitioning on HP servers refers to the ability to run multiple instances of the HP-UX operating system on a single server. Partitioning is accomplished by allocating different subsets of the server's hardware resources to different [partitions](#), each of which can run a single instance of HP-UX, with isolation from other partitions.

There are two main types of partitions that can be created on HP servers: hardware partitions (*nPars*), and virtual partitions (*vPars*). Depending on the server model, one or both types of partitions may be supported. The current version of Partition Manager can create, delete, and

administer only *hardware* partitions. For this reason, the unqualified term *partition* refers specifically to hardware partitions wherever it is used in the Partition Manager program and this help system. The terms *virtual partition* and *vPar* are used when referring to virtual partitions.

hp Partition Manager is used for partition management.

For additional sources of information about partitioning of HP servers, Partition Manager documentation which can be found here:

<http://info.fc.hp.com/projects/parmgr/11.11/help/C/overview.html>.

Hardware Partitioned Systems

A *complex* consists of a single hardware configuration that can support multiple instances of an operating system, by means of hardware partitions.

A partition is a configuration of *cells* (processors and memory), and one or more *I/O chassis*. An I/O chassis has slots that may contain I/O cards that connect the partition to peripheral devices.

A complex comprises one or more cabinets that are cabled together. Single-cabinet complexes include the HP server rp8400 (model number *S16K-A*), HP server *rp7410*, HP Superdome *SD16000*, and HP Superdome *SD32000*. Multiple-cabinet complexes include the HP Superdome *SD64000*, which is formed by cabling together the crossbars (*XBCs*) from two adjacent SD32000-type cabinets. Additional systems supporting hardware partitions may be provided in the future.

Information about the configuration of a complex is stored in the *Complex Profile* maintained by the *service processor*. Several Partition Manager tasks require the service processor to read and update the Complex Profile.

Partitioning and Multiple Windchill instances

A complete Windchill environment may require separate Windchill instances for various distinct projects as well as Production, Pre-Production, Training, Integration and Development platforms. To consolidate cost and system management, customers can use the partitioning described above to put multiple Windchill instances on the same machine divided by vPars or nPars. The memory and CPU requirements below are minimums for each individual Windchill instance regardless of the platform (Production, Pre-Production, etc.) targeted. This way, one has the choice of, for example splitting the implementation into partitions for each project or each platform. For example, if we have Project A and B for which we want Production, Pre-Production and Training/Integration/Development environments, we can use an rp8400 divided into 4 vPars of 4-CPU each with Project A and Project B Production environments on the first partition, Pre-Production environments on the second partition and Training/Integration/Development environments on the third partition, and the Oracle database on the fourth partition. In this case, we probably want 8 GB RAM on each partition - enough to account for both Windchill projects depending on their final user counts. Please see your local HP Services Consulting and Integration (HPS C&I) consultant for more information and sizing help.

Also note that a system is recommended to contain at least 2 GB per CPU. The guidelines below are specifically for calculating how much memory/CPU is required for multiple Windchill instances but that the bare minimum is 2 GB per CPU as mentioned above.

Table 2 - Memory/CPU for Windchill partitions

User Load	3-Tier Windchill Partition		3-Tier Oracle Partition		Monolithic Partition	
	Number of CPUs	Memory in GB	Number of CPUs	Memory in GB	Number of CPUs	Memory in GB
100	2	2	1	1	2	4
500	2	3	2	2	3	5
1000	3	5	2	2	5	7
2000	6	6	4	2	8	6
5000	16	17	6	2	20	19

Choosing Disks

HP has technical consulting available in choosing disk configurations appropriate to certain architectures. What follows is a quick overview of some of the options to consider. Also note that the key measurements used to determine performance of disks are IOPs/second and throughput in MB/second.

Table 3 - Disk options

Requirement	ds2100	ds2405	ds2300	ra3000-rm	va7110	va7410	xp128 / xp1024
Built-in Disk array management and security				X	X	X	X
Fiber Channel connect					X	X	X
Rack-Mount	X	X	X	X	X	X	
Mixed environment of NT and HP-UX				X	X	X	X
Performance but no security	X	X	X				
JBOD and Mirror/UX	X		X		X		
Capacity							
Less than 1Tb	X						
1Tb and above			X		X	X	X

Configuring for High Availability

Mirror/UX

Mirror/UX offers some rudimentary LVM RAID1 mirroring capability that is interesting if you do not have an AutoRAID or other disk array, like the HP FC60 or XP256 or XP512. It is a very fast

mechanism and is very reliable. During the scalability testing, a 17 Gb mirror was demonstrated to be synchronized with an offline copy in about 40 seconds.

MC/ServiceGuard

MC/ServiceGuard is an optional software product that enables the creation of multi-node server clusters designed to ensure application availability. MC/ServiceGuard uses the concept of an application 'package' to define the service(s) running on a node within the cluster. In the event of a hardware or application failure, the package(s) will be migrated to a surviving node. Package definition files and associated scripts have been created for Windchill, Oracle, and Netscape Enterprise Server. These scripts are available from the www.hp.com/go/pdm site in the Tech Support area. Additionally, a Windchill & MC/ServiceGuard configuration white paper may be found on the same web site.

MC/ServiceGuard does *not* require dedicated, 'hot-standby' hardware to achieve High Availability. All nodes in a cluster may be active and running separate packages. When a failure is detected, the cluster's workload may be redistributed across the remaining nodes.

Because of the cluster IP addressing involved in the MC/ServiceGuard configuration, there are several things to take into account in the configuration and installation of Windchill (assuming that you are installing Aphelion/Windchill on the first server and RetrievalWare/Oracle on the second machine) :

- 1/ When installing all the products, ensure that the hostnames always refer to the cluster server names and not the primary local addresses on the machine.
- 2/ In Info*Engine, the Windchill adapter will fail to start. This is avoided by adding the following property through the LDAP Browser for the Windchill adapter in Info*Engine:
com.mycompany.Windchill.socketAccess.clusterName=cluster.mycompany.com
- 3/ For RetrievalWare, you will need to duplicate the HOST and ENVIRONMENT ON_HOST clauses in \$RWARE_HOME/admin/exec.cfg and \$WINDCHILL_INDEXES/config/exec.cfg so that both the cluster names and the physical names of the machines are listed. Additionally, you will need to specify a -host option to the execd command line.
- 4/ For Aphelion LDAP Directory, you will need to install it on both systems and remove the binaries on the second system (the place where you installed Aphelion). This will leave the links in /opt, /usr, and /var that are necessary for Aphelion to start up in the case of a failover to this second system.
- 5/ Try and install everything into a minimum number of mounts point for ease of packaging. For example, you can install Apache/Tomcat/Windchill/Aphelion in a directory like /wt626, Oracle into /opt/oracle and Apache/Tomcat/RetrievalWare/WindchillSearch into a directory like /rware.

More information on MC/ServiceGuard may be found at www.hp.com/go/ha.

Configuring for Security

Reverse Proxy Server

Reverse proxy servers can be used for limiting access to the Windchill application servers from external clients. Client requests are received by the reverse proxy server residing on the external network or DMZ and are routed to the internal, firewall protected, application server. The

external clients use a site url and never see the internal network addresses. Increased security is possible using two reverse proxy servers on a dual homed host, with each server reading and writing to only one of the network interfaces for external and internal communication with data being exchanged between the servers. This is used in conjunction with a secure or hardened operating system such as the Virtual Vault OS and forthcoming products from HP.

Firewalling

For a detailed report on firewalling and Windchill, please see Appendix B of the Windchill System Administrator's Guide.

Planning a Backup Strategy

Omniback/Data Protector

Omniback/Data Protector can be implemented easily to perform offline backups. Its interface is self-explanatory and the only thing to remember is to shutdown the database and Windchill before starting the backup. If you are using a file vault, do not forget to backup the database and the vault at the same time. The one drawback is that if you have a failure, your only protection is this complete backup. This means that you will lose everything up to the last backup and will have to restore the entire backup image to get back online.

Oracle Archiving and Hot Backup

If you are exclusively using BLOBS for data storage and if you enable archive logs in Oracle, you can perform a backup online without impacting the connected users. Oracle will continue to write to the redo logs while Omniback backs up the database files. Once the backup is complete, Oracle rolls forward the transactions stored in the redo logs and will continue working normally. Once again, if a vault is used, it needs to be integrated into the overall strategy. One idea for vault backup is that the vault is mirrored to begin with; before the backup the mirror is split. One copy continues online getting new transaction changes. The split copy is backed-up with the database and then the old copy is synchronized and brought up-to-date with the new copy.